



PCT/AU00/01023

REC'D 19 SEP 2000

WIPO

PCT

**PRIORITY
DOCUMENT**

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

**Patent Office
Canberra**

AU 00/01023

EU

I, KAY WARD, ACTING MANAGER EXAMINATION SUPPORT AND
SALES hereby certify that annexed is a true copy of the Complete specification
in connection with Application No. 47400/99 for a patent by ALCATEL filed on
06 September 1999.

10/070311



WITNESS my hand this
Thirteenth day of September 2000

K Ward

KAY WARD
ACTING MANAGER EXAMINATION
SUPPORT AND SALES

AUSTRALIA

Patents Act 1990

ORIGINAL
COMPLETE SPECIFICATION
STANDARD PATENT

Invention Title:

“RECURSIVE TRAFFIC DISTRIBUTION IP/DATA NETWORK MODEL”

The following statement is a full description of
this invention, including the best method of
performing it known to us:

Technical field

This invention relates to a method and arrangement for transferring network node status information between nodes.

Background Art

5 Least cost routing examines the network topography to determine the shortest path between the source of a message and its destination. This method does not take account of the load status of the individual links of the chosen path so one or more of the links may become overloaded, preventing or disrupting the delivery of the message to the destination.

10 An alternative proposal is to take into account the load status of the links when determining the chosen path. Thus the system may determine all the shortest paths and make the path selection on the basis of the path with the links carrying the least traffic. To implement this, it is necessary for all the nodes to exchange load status information. As the number of nodes and links in a network grow, the
15 implementation of this technique requires the exchange of a large amount of load status traffic, as each node broadcasts the load status of its associated links to the other nodes.

 In our co-pending application number 44470/99 (127045 SY), we disclose a network in which the message is spread over all the practical paths between the
20 source and destination. An advantageous embodiment of that invention involves the elimination of paths having heavily loaded links. Again, the implementation of this technique requires the exchange of load status information between nodes, generating a large volume of traffic.

Disclosure of the Invention

25 This specification discloses a network arrangement for a plurality of nodes each node being connected to one or more other nodes by corresponding node links,

- the network being arranged into a recursive hierarchy of units having two or more levels,
- 30 - the nodes being the units of the first level of the hierarchy,

- the units of higher levels of the hierarchy being formed by groupings of the units of the previous level,
- wherein the units of a level exchange a corresponding load status information.

5

Brief Description of the Drawings

Figure 1 is a schematic representation of a network in which the nodes are arranged in a recursive hierarchy, in accordance with an embodiment of the invention.

10 Figure 2 represents a load status monitor.

Figure 3 illustrates the message structure for exchanging information at different levels.

Best Mode of Carrying out the Invention

15 Figure 1 shows a network of nodes interconnected by links. According to the embodiment shown in Figure 1, the nodes are linked in a logical hierarchy.

~~The nodes exemplified by the circles, some of which are numbered 101 ... 202 ... 303 ..., are shown interconnected by node links, represented by the lines drawn between the nodes.~~

20 The nodes are formed into groups 10, 11, 12, 20, 30. The groups are interconnected by group links, for example 1001 between node 105 of group 10 and node 121 of group 12.

The group links are shown in the following Table 1.

TABLE 1 : GROUP LINKS

Node/Group		Node/Group		Link
105	10	121	12	1001
112	11	122	12	1002
103	10	111	11	1003
123	12	202	20	1201
101	10	302	30	1301
102	10	303	30	1302

Preferably the group links have a larger traffic capacity than node links. Group links may equate to regional trunks within a particular carrier's network, or to links between different carriers, different countries or different global regions, for example.

5 As shown in Figure 1, the levels of the hierarchy start with the nodes. The units of the second level are the groups of nodes. The units L31, L32, L33, of the third level are one or more groups, and the units of the fourth level, L41, L42 are formed by aggregating units of the third level. Thus L41 is the aggregation of L31 and L33, while L42 encompasses L32.

10 Preferably the nodes are grouped on the basis of communication path topography so that there are relatively few links in the shortest path between any two nodes in a group.

The groups themselves, 10, 11, 12, 20, 30 are linked by designated links 1001, 1002, 1003 1201, 1301, 1302 connecting designated nodes in the
15 corresponding groups.

Within each group, a master node is assigned, 104, 113, 124, 201, 301. Because each node in a group knows the load status of all the links in that group, the master node has the information to enable it to compile available capacity reports adapted to meet the requirements of the higher levels.

20 The physical node interconnections are illustrated at A in Figure 1 and B,C and D illustrate the conceptual logical links for the second, third and fourth levels of hierarchy.

At level A, the nodes of each group communicate load status information to each of the other nodes of the corresponding group.

25 Thus the master nodes 104, 113, 124, of the network 1 shown at B are aggregated under the supervision of a single master node 113 which is assigned to the next higher level. Because, in the example shown, networks 2 and 3 have only one group, the same master 201, 301 is used all levels.

30 At level C, the master nodes 113, 201, 301 exchange information as to the overall load status of their associated networks 1, 2, and 3.

The nodes at level C are then 113, 201, 301. By then aggregating 113 and 301, one of these two can be designated to manage the level D information exchange for both network 1 and network 2. Thus at level D, 201 and 301 can exchange information on the available capacity for the regions covered by L42 and L41.

At level C, information on the capacity of networks 1, 2 and 3 is exchanged between the networks. At level B, information on the capacity of groups 10, 11, 12, 20, 30 is exchanged between the groups.

At level A, the nodes interchange information on traffic capacity at the node link level, within the groups.

Figure 2 shows an arrangement for monitoring the available capacity of the links connected to a node.

For each link connected to a node there is a buffer 51, 52, 53, e.g. in the form of FIFO.

The traffic level monitor 50 checks the level of the contents of the buffers to measure the available capacity on the basis of the speed of the link associated with the buffer. The result of the monitoring is then reported to the other nodes in the same group.

In a simplified measuring system the monitor may report whether or not a link has spare capacity, e.g. by checking whether a buffer's content is above or below a predetermined threshold.

The load status information exchange is carried out on the following basis.

The nodes within a group each notify the other nodes within that group of the load status of the links connected to the notifying node.

At the group level, each group notifies the other groups of the load status of the links connected to the notifying group and of the load status of internal paths within the group available for interconnecting the group links connected to the notifying group.

For example, Group 12 is connected to Group 20 via link 1201, to Group 10 via link 1001, and to Group 11 via link 1002.

Preferably, the designated as a master node manages the interchange of information between the groups.

Table 2 shows the master nodes for each group.

TABLE 2

5

GROUP	MASTER NODE
10	104
11	113
12	124
20	201
30	301

As can be seen in Figure 1, the master nodes take part in the higher level exchanges but their number is progressively reduced by the recursive grouping.

Thus, in the embodiment shown in Figure 1, while there are 5 master nodes
10 shown at level B, there are only 3 at level C and 2 at level D.

Preferably, the grouping is carried out on the basis of proximity in the sense of the number of links in the path. Of course this is not a strict rule at the node level because the nodes at either end of a group link are joined by a single link, while there may be more than 2 links between nodes within a group. Other factors which
15 influence grouping are geographical proximity and network ownership, as well as the traffic flows.

For example, the nodes of network 2 may be geographically close to node of network 1, but network 1 may be owned by a different carrier from network 2.

At level D, the nodes 201 and 301 exchange information on the available
20 capacity between network 3 and network 2. This information would, for example, be based on the load status of links 1201, 1301, 1302, and the capacity across network 1 between link 1201 and the links 1301, 1302. The information need only identify the maximum available capacity at the time, which varies in accordance with the load on the various network elements.

25 For the sake of clarity the information will be given the following names:

Level D = Regional;

Level C = Network;

Level B = Group;

Level A = Node.

5 Regional information may be, for example, the maximum available capacity between the "electrically" remotest groups. The term "electrically" refers to the number of links and may include cable, optical and radio links.

10 Network information may be, for example, the capacity between the various networks, including the trans-network capacity between the network links 1201, 1301, 1302.

 Group information could be typified by the capacity between groups, including the trans-group capacity between the group links.

 Node information is the information broadcast by a node to the other nodes within its group as the load status of the node and its associated links.

15 Group information can be deduced from node information. Each node in a group knows the load status of all the nodes in that group. Thus the master node 124 in group 12 knows the status of group links 1001 from node 121, group link 1002 from node 122, and group/network link 1201 from node 123, as well as the status of all the internal nodes and links within group 12. Node 124 can therefore

20 calculate the available capacity across the group 12 between any pair of the links 1201, 1001, 1002. Preferably the master node 124 would use the "all practical paths" algorithm of our Australian Patent application 44470/99 (Docket No. 127045 SY) to calculate the trans-group capacity. This group information is

 interchanged between the group master nodes 201, 124, 113, 104, 301 at level B.

25 The units of the level B group domain are again grouped together, in this embodiment, into 3 network groups. The network groups include two one member groups 201 and 301, and one three member group 124, 113, 104. The network master of each one member group is the member of the group, while 113 is designated as the master of the three member network group.

30 The three network masters from level B interchange network information at the level C network domain. The information relates to the network links connecting

the respective networks, and the trans-network information relating to the capacity between the pairs of network links.

At the regional domain, level D, the network masters 201, 113, 301 have been formed into two groups, resulting in two regional masters 201, 301, which
5 exchange information on the available capacity between the two regions.

The regional master nodes 201, 301, convey the regional link capacity information to the other regional nodes. In the present embodiment 301 conveys the information to 113. 201 is the only regional node in the other regional grouping.

The regional nodes 201, 113 and 301 are all network master nodes and they
10 convey the inter-regional and inter-network capacity information to the network level nodes. In our embodiments, 113 conveys this information to the nodes 104, 124.

Each of the network level nodes 201, 124, 113, 104, 301 is a group master and relays the higher level information to each of the nodes in its group.

The grouping of the units at each level means that the information exchanged
15 at each level becomes more generalised.

This means that a node has detailed capacity information about the other nodes in its group. Capacity information about other groups in its network, capacity information about the other networks in its region, and information about the inter-regional capacity.

20 In a preferred embodiment the group master handles the interchange of node link capacity information. Each node, instead of broadcasting its load status to all the other nodes in the group, sends the information only to the group master, which collates the information from each node and relays the information to the other nodes. The message from the group master preferably incorporates the higher
25 level load status information, so that each node has an overall picture of the entire system.

Thus the group master may broadcast a message including the information shown in Figure 3. The first segment RL includes the load status at the regional link level D. A second portion of the payload includes a number of segments of
30 information on the inter-network load status NL. A third portion includes segments

GL on the inter-group load status, and the fourth portion includes segments NL on the load status of the nodes within the group.

The claims defining the invention are as follows:

1. A network arrangement for a plurality of nodes each node being connected to one or more other nodes by corresponding node links,
 - the network being arranged into a recursive hierarchy of units having
 - 5 two or more levels,
 - the nodes being the units of the first level of the hierarchy,
 - the units of higher levels of the hierarchy being formed by groupings of the units of the previous level,
 - wherein the units of a level exchange a corresponding load status
 - 10 information.
 2. An arrangement as claimed in claim 1 wherein, within, each group of units, a master entity is designated, the master entity conveying inter-unit load status information relating to the units of that level to the next higher level.
 3. An arrangement as claimed claim 1 or claim 2 wherein, in the first level, a
 - 15 selected node in each group is designated as the master node for the corresponding group,
 - the master node managing the transfer of node load status information within its corresponding group.
 4. An arrangement as claimed in claim 1 or claim 2, or claim 3 wherein the
 - 20 load status information includes information on the available traffic capacity between the ports of each unit.
 5. An arrangement as claimed in any one of claims 1 to 4 wherein each node includes node load status monitoring means to monitor the load status of the links connected to the node.
 - 25 6. An arrangement as claimed in any one of claims 1 to 5 wherein at least one node of each second level group is connected to a node of at least one other second level group via a corresponding group link whereby group load status information can be interchange.
 7. An arrangement as claimed in claim 6 wherein the units of the third level are
 - 30 formed by mutually interconnected second level units.
-

8. A network arrangement for interchanging load status information substantially as herein described with reference to the accompanying drawings.

5

DATED THIS EIGHTEENTH DAY OF AUGUST 1999
ALCATEL




FIGURE 1




FIGURE 2

RL	NL	NL	NL	GL	GL	GL	GL	GL	GL	NL	NL	NL	NL	NL	NL	NL
----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

FIGURE 3

ABSTRACT

To avoid the need for all the nodes of a network to know the load status of all other nodes and links in the network, the nodes 101 303 are organized in node groups 10, 11, 12, 20, 30. Each group has a master node, 104, 124, 113, 201, 301 which is incorporated in a logical higher order group. By recursively grouping the units of each preceding group, a logical hierarchical structure is formed in which the number of units at each level A, B, C, D, decreases.

Each node informs the other nodes of its load status, and each higher level unit also exchanges information with the other units. The recursive, hierarchical structure greatly reduces the amount of load status information exchanged across the network.
